



September 12, 2014

Via Electronic Filing

Marlene H. Dortch, Secretary
Federal Communications Commission
445 Twelfth Street, SW
Washington, DC 20554

Re: *Written ex parte presentation, GN Docket No. 14-28*

Dear Ms. Dortch:

Mobile Future respectfully submits to the Commission the attached paper, “How Wireless is Different: Considerations for the Open Internet Rulemaking” written by engineer Peter Rysavy of Rysavy Research. The report uses a data-driven and scientific engineering perspective to address key differences in wireless networks when compared to wireline, making the case for flexible, nimble network management practices and operations. Rysavy notes that, “subjecting wireless broadband networks to rules that dictate how wired broadband networks are designed and operated would be a mistake.”

Despite historic levels of capital investment to address increasing consumer demand, including expanding and upgrading networks, harnessing small cells and offloading onto Wi-Fi, wireless networks remain inherently capacity-constrained when compared to fixed broadband networks. As an example, the report notes, “all usable radio frequency spectrum, across thousands of licensees, has less than 10 percent of the capacity of a single fiber-optic cable.”

Rysavy concludes that given the special characteristics of wireless networks, and the implicit spectrum capacity challenges in which mobile networks operate, new regulatory mandates to treat all data packets equally will impede the management of mobile traffic and congestion. Given the surge in demand for wireless services, such mandates would be counterintuitive, will not serve the public interest, and may derail future innovation.

Pursuant to Section 1.1206 of the Commission's rules, a copy of this letter is being filed via ECFS with your office. Please do not hesitate to contact the undersigned with any questions.

Respectfully submitted,

/s/ Jonathan Spalter

Jonathan Spalter, Chair
Mobile Future
1325 Pennsylvania Avenue, N.W., Suite 600
Washington, D.C. 20004
202-756-4154
www.mobilefuture.org

Attachment

cc: Ruth Milkman
Gigi Sohn
Renee Gregory
Adonis Hoffman
Louis Peraertz
David Goldman
Matthew Berry
Brendan Carr
Erin McGrath
Julius Knapp
Roger Sherman
Jim Schlichting
John Leibovitz
Jessica Almond



How Wireless is Different

Considerations for the Open Internet Rulemaking

September 12, 2014

Table of Contents

NOTICE	2
EXECUTIVE SUMMARY	3
INTRODUCTION	4
CAPACITY	5
SPECTRUM	6
EFFICIENCY	7
NUMBER OF CELLS	8
CAPACITY ANALYSIS	9
WIRELESS NETWORK TRAFFIC MANAGEMENT.....	11
VOICE, VIDEO, AND QUALITY OF SERVICE.....	13
CONCLUSION	14
APPENDIX.....	16
ABOUT RYSAVY RESEARCH	22

Notice

Rysavy Research provides this document and the information contained herein to you for informational purposes only. Rysavy Research provides this information solely on the basis that you will take responsibility for making your own assessments of the information. Although Rysavy Research has exercised reasonable care in providing this information to you, Rysavy Research does not warrant that the information is error-free. Rysavy Research disclaims and in no event shall be liable for any losses or damages of any kind, whether direct, indirect, incidental, consequential, or punitive arising out of or in any way related to the use of the information.

Executive Summary

At no point in the evolution of American innovation policy has there been a greater need for lawmakers and regulators to understand how wireless broadband communications networks differ from wireline. Policies arising from overlooking this difference could have multiple and prolonged negative consequences to mobile consumers and our economy. Despite any indication of real-world harm, some are renewing calls for additional network neutrality rules—and urging the Federal Communications Commission to subject both wireline and wireless networks to this new regime.

Subjecting wireless broadband networks to rules that dictate how wired broadband networks are designed and operated would be a mistake given the inherent capacity limitations of radio spectrum and the resulting differences in how mobile networks are designed and managed. While mobile networks can deliver a broadband experience that rivals some forms of wireline systems, often in the tens of megabits per second, these networks are nevertheless too capacity-constrained due to their inherently unique technical and operational characteristics to consistently offer that level of service to all subscribers at the same time.

Wireless capacity has remained only slightly ahead of rapidly rising mobile broadband demand despite massive investments in network infrastructure and consistently increasing use of more spectrum. This paper's intent is not to suggest that imposing network neutrality restrictions on wireline networks is necessary or easy or that various forms of network management are not important to providing high quality wireline services. Instead, this paper explains the unavoidable consequences of imposing rigid rules on capacity-constrained wireless networks.

Limiting or directing by government fiat the ways in which wireless network operators can respond to surges in demand for more capacity will undermine the user experience, impair quality of service, and limit accessibility of hundreds of millions of wireless subscribers at precisely the same time policymakers are trying to entice more Americans, including those living in underserved and rural communities, onto the country's digital highways. As importantly, such restrictions will delay the deployment of new next-generation technology, undermining the investment and innovation that are driving the mobile revolution and apps economy.

Three factors determine capacity: the amount of available spectrum, spectral efficiency, and number of cell sites. The wireless broadband community is dedicating enormous capital to address the capacity challenge from all three dimensions. Ultimately, however, wireless networks have only a fraction of the capacity available to wireline networks at any point in time. To convey a sense of the scale of the chasm: All usable radio frequency spectrum, across thousands of licensees, has less than 10 percent of the capacity of a single fiber-optic cable. As a consequence, just a small percentage of the subscriber base in a coverage area streaming video can completely consume local capacity.

In addition, due to the way radio signals propagate through environments, the quality of a radio signal can vary tremendously. Also variable is the number of users at any one time in a coverage area—the people using these networks are, by definition, mobile. Consequently, wireless operators must carefully, constantly, and flexibly manage their networks using multiple mechanisms to optimize precious radio resources.

Operators also must dynamically prioritize among types of services, for example giving voice higher priority than data. Voice over LTE, in particular, depends on a quality-of-service (QoS) architecture in LTE so that voice packets arrive uninterrupted. QoS could also play an important role for other applications, including public safety, health care, the smart grid, and education. Rules that restrict QoS management will eliminate a potentially effective way to ensure wireless networks better serve consumers and the public interest, as well as conceivably eliminating entire classes of applications that haven't even been invented. A regulatory mandate to treat all packets and circumstances equally in a capacity-constrained environment does not make engineering sense, nor serve the public interest, and will almost certainly derail future transformative innovation while degrading U.S. consumers' mobile experience.

Introduction

Mobile broadband networks differ from wireline networks in a number of fundamental ways, including the inherent capacity limitations of radio spectrum and the related implications on wireless network design and operational management. Mobile broadband technology also remains in a nascent stage, as broadband providers race to deploy 4G LTE service while already turning an eye to next-generation technology, such as 5G, that will further enhance service and allow more direct competition with wireline offerings.

The differences between wireless and wireline must be properly considered in the FCC's current proposed rulemaking, "Protecting and Promoting the Open Internet." This paper explains the profound differences, including capacity and options to augment capacity, how wireless networks must be managed at multiple layers, the dynamic nature of wireless connections, and how different applications have diverse requirements when operating over wireless networks.

The wireless networks discussed in this paper are those using mobile technologies such as Long Term Evolution (LTE), not Wi-Fi networks. Wi-Fi is a local-area technology and has different capabilities and constraints.

Given the continued astounding success of mobile broadband, regulators should exercise extreme caution before imposing new constraints on these networks and their operations that could undermine the full, untapped potential of mobile broadband services.

In particular, given the significant difference between wireline and wireless broadband, rules intended for wireline networks could have unintended and negative consequences if applied to mobile broadband

networks. For example, voice over LTE depends on a quality-of-service architecture that prioritizes voice traffic over data traffic. If that prioritization were forbidden, 4G voice would be rendered unreliable and in many cases unusable because voice packets would have to compete with data packets.

The wireless industry is innovating and changing at a rapid rate. Its success is, in part, a result of light regulation. The United States is now the world leader in mobile broadband technology development and deployment. The smartphone revolution began with the launch of the iPhone and the development of Android, both of which occurred after the 2007 order classifying mobile broadband as an information service. Similarly, the 4G revolution occurred after the FCC's 2010 "light-touch" network neutrality rules that treated wireless differently than wireline. New rules that allegedly protect this dynamic ecosystem but that in fact impose new restrictions are unlikely to provide any benefit and may well produce harmful and unintended consequences.

Capacity

As demand for mobile broadband continues to grow unabated, the wireless industry has exploited every avenue to augment capacity, including deploying more efficient technology, using more spectrum, adding more sites, harnessing small cells, and offloading onto Wi-Fi.

The gargantuan upgrade from 3G to 4G LTE, touching every cell site, provided huge gains in efficiency. Now LTE-Advanced, in its earliest stages of deployment, achieves even higher efficiency with smart antennas and the ability to aggregate radio carriers.

Wireless capacity has remained only slightly ahead of demand, depending on constant investments in network infrastructure and use of more spectrum. These efforts have enabled wireless networks to maintain pace with demand, but even with advances such as 4G, the capacity of wireless networks remains far more constrained than that of wireline networks.


From a consumer perspective, the effects of exceeding capacity are dramatic. Once the network is operating close to capacity, throughputs go down, delays go up, and applications start performing sluggishly. As congestion increases, applications can timeout or fail entirely. Erratic and unreliable operation not only thwarts use but discourages developers and innovators, including those developing "over-the-top" applications that provide voice and multimedia capability using data connections.

Any broadband network consists of a core network built on fiber and a "last-mile" technology. That last-mile technology can be wire, such as coaxial cable or DSL; fiber; or radio frequencies. The benefit of radio is mobility, which appeals to people since they are inherently mobile, but the tradeoff is lower capacity, particularly in mobile networks that serve large areas.

As shown in Figure 1, three dimensions determine the capacity of a wireless network: the amount of available spectrum, the spectral efficiency of the technology, and the number of cells in a network. To

understand why wireless network capacity cannot be easily increased requires an understanding of the constraints affecting each of these three dimensions, including how they differ from wireline networks.

Figure 1: Three Dimensions Determining Wireless Network Capacity



Spectrum

Spectrum translates directly to capacity. Spectrum, however, is a limited and finite resource for mobile broadband networks. Radio spectrum extends from about 3 kilohertz (KHz) to about 300 gigahertz (GHz), but current technology limits mobile use to frequencies under 3 GHz. Of those, U.S. cellular systems use roughly 500 MHz. An individual operator's access to spectrum is much smaller, varies by market, and is subject to spectrum aggregation rules.

In contrast, wireline networks have access to far more frequencies in the mediums they use, and these frequencies can be compared with radio spectrum. Not only do coax cable and especially fiber-optic cable have access to a vast array of frequencies, they also carry their spectrum within a controlled physical medium over which the operator has near-complete control. As soon as the capacity of one fiber-optic cable is exhausted, another one is available alongside. In contrast, wireless networks rely on propagation through the air, and the same frequencies cannot be used without interference until some distance away.

A goal of 5G is to harness higher frequencies, such as from 10 GHz to 100 GHz. Challenges at these frequencies are daunting, but likely surmountable for some scenarios, such as small cells. However, 5G networks, completely undefined at this time, are not likely before 2020, so any capacity relief from use of higher frequencies will not occur until such networks are widely deployed, unlikely before the 2025 timeframe.

In the meantime, forthcoming mobile broadband spectrum from sources such as the television incentive auctions and Advanced Wireless Services-3 (AWS-3), while hugely welcome and necessary, will not fundamentally change the inherent limits of radio spectrum capacity.

Efficiency

Given limited spectrum and huge demand, wireless networks must extract as much capacity as possible from available spectrum. Today's networks, using sophisticated modulation and encoding methods, already operate at close to maximum theoretical spectral efficiency. These methods are constrained by the laws of physics, including the Shannon bound, which dictates the maximum achievable spectral efficiency as measured by bits per second (bps) of throughput per Hertz (Hz) of spectrum relative to the amount of noise.¹

Unlike computing, in which better-designed chips translate to continually faster speeds, increasing efficiency in radio technology is far more challenging. Engineers will keep increasing efficiency, but no magic exists to suddenly increase capacity by orders of magnitude. Instead, gains are incremental—10 percent here, 20 percent there.

The most promising avenue on the wireless technology roadmap is employing more advanced smart antennas, but this option requires upgrades to every single base station to install additional antennas.

For example, operators are currently upgrading antennas and radios in LTE networks from a 2X2 Multiple Input Multiple Output (MIMO) downlink configuration (two antennas at the base station, two at the mobile) to a 4X2 configuration (four antennas at the base station, two at the mobile), an expensive upgrade that will yield approximately a 20 percent efficiency gain.²

The concept of efficiency also applies to wireline networks, in which achieving the fastest possible speed from a digital subscriber line or coaxial cable also requires advanced modulation and coding. A radio connection, however, is inherently more variable and challenging than a wire or fiber cable, so the modulation and coding must be dynamic in order to adapt to instantaneous radio fluctuations.

¹ For background on spectral efficiency, see Peter Rysavy, "Challenges and Considerations in Defining Spectrum Efficiency," *Proceedings of the IEEE*, Volume 102, Issue 3, 2014. Available at <http://www.rysavy.com/Articles/2014-03-IEEE-Defining-Spectrum-Efficiency.pdf>.

² Rysavy Research, "Beyond LTE: Enabling The Mobile Broadband Explosion," *4G Americas*, August, 2014, page 71. Available at <http://www.rysavy.com/Articles/2014-08-4G-Americas-Mobile-Broadband-Explosion.pdf>.

Radio connections also have more noise to contend with than shielded wires, limiting modulation methods. For example, the latest version of Data Over Cable Service Interface Specification (DOCSIS) permits modulation of up to 4096 Quadrature Amplitude Modulation (QAM), much higher than the 64 QAM possible with LTE. With 4096 QAM, each symbol (waveform modification) encodes 14 bits of data, whereas 64 QAM encodes only 6 bits of data. Even with 64 QAM available, only mobile devices close to the base station can take advantage of this higher-order modulation.

Number of Cells

The number of cells also determines the capacity of a network. Assuming each cell has a certain amount of capacity, increasing the number of cells by a factor of N increases capacity by that same factor. Historically, the number of cell sites has contributed more to capacity than any other factor. In 1993, the United States had 12,824 sites. Twenty years later, it had 304,360 sites, an increase by a factor of 24.³ Although an impressive gain, the increased number of sites has not allowed capacity to come even close to matching the capacity of wireline networks.

In developed countries, macro cells already cover population areas almost completely. The next wave of cellular deployment is for capacity—densifying networks. The smaller the cell, the greater the capacity per subscriber; this is why Wi-Fi access points, with their small coverage areas, can offer such high throughputs per subscriber.

Small cells, such as femtos in homes and indoor and outdoor picocells, can yield large capacity gains, but multiple difficulties must be addressed, including backhaul, power, management, interference, site acquisition, and local zoning approval. Although one can imagine a future with many millions of small cells, getting to that future will be a multi-year process, one not likely to be realized until the end of this decade or later.

The concept of cells does not directly apply to wireline networks, but analogies can be made. For example, in cable systems, the operator extends fiber to nodes, from which emanate separate coaxial cables, each serving a number of homes or buildings, which in aggregate can be thought of as a coverage area. The television stations on that cable are broadcast, while the Internet spectrum and capacity are shared by a certain number of users, similar to what happens in wireless networks.

What differs from wireless is that the coax is strung to specific buildings as needed, whereas a wireless operator has to provide continuous coverage over a much greater number of square

³ CTIA, *Semi-Annual Wireless Industry Survey*, 2013. CTIA, “Annual Wireless Summary Report, Year-End 2013,” 2014, <http://www.ctia.org/your-wireless-life/how-wireless-works/annual-wireless-industry-survey>, accessed July 17, 2014.

kilometers. Additionally, installing base stations with towers and antennas, along with connections to those base stations, represents extremely expensive infrastructure.

Capacity Analysis

Table 1 summarizes the capacity factors for wireline and wireless broadband networks.

Table 1: Summary of Capacity Factors Between Wireline and Wireless

Factor	Wireline	Mobile Wireless
Spectrum	Huge amounts of spectrum available, especially with fiber-optic cable.	Radio spectrum capacity inherently constrained.
Spectral Efficiency	Extremely high. Modulation and coding operates in a stable, predictable, low-noise environment.	High and must be able to adapt to a continuously varying communications environment.
Number of Cells	Concept does not directly apply.	Huge historic increase. Next wave will be small cells, but these have to overcome difficult challenges.

The resulting differences in capacity are profound. Even employing the capacity of all radio frequency to 100 GHz results in less than 10 percent the capacity of a single fiber-optic cable.⁴

Table 2 shows different LTE configuration options, bandwidth requirements for some popular types of applications, the number of simultaneous users the network can support in any cell sector, and the percentage of subscribers this number represents assuming three hundred and fifty subscribers for that cell sector.⁵

⁴ One fiber optic cable can transmit over 10,000 Gbps, compared to all wireless spectrum to 100 GHz that even at an extremely high 10 bps/Hz would have only 1,000 Gbps of capacity.

⁵ Year-end 2013, the U.S. had 335.65M subscribers and 304,360 cell sites, resulting in an average of 1103 subscribers per cell and thus 1103/3 or 368 subscribers per sector for three-sector cells. CTIA, Annual Wireless Summary Report, Year-End 2013.

Table 2: LTE Configurations and Number of Users Supported for Different Applications⁶

	Bandwidth (Mbps)	LTE in 5 MHz + 5 MHz	LTE in 10 MHz + 10 MHz	LTE in 30 MHz + 30 MHz
Downlink Sector Capacity (Mbps)		8.5	17	51
YouTube at 480p resolution	1			
Number of simultaneous users		9	17	51
Percentage of subscribers		2%	5%	15%
Netflix at 720p resolution	3			
Number of simultaneous users		3	6	17
Percentage of subscribers		1%	2%	5%
Netflix at 1040p resolution	5			
Number of simultaneous users		2	3	10
Percentage of subscribers		0%	1%	3%
Netflix at 4K resolution	15			
Number of simultaneous users		1	1	3
Percentage of subscribers		0%	0%	1%

The table quantifies just how constrained capacity is relative to the applications that users enjoy. Even LTE networks in an extremely high-capacity configuration using 60 MHz of spectrum can support only a small percentage of a subscriber base consuming high-bandwidth applications.

Whereas wireline broadband networks can deliver streaming video to large numbers of users, just 5 percent of subscribers streaming YouTube in a typical LTE cell sector today can consume the total capacity of that sector. This type of usage is critical context for policymakers given the fact that mobile devices now account for 21 percent of all online video traffic—having grown at a compound annual growth rate of 152 percent since 2012.⁷

⁶ Initial LTE deployments, typically using 700 MHz spectrum, are either in configurations of 5 MHz downlink + 5 MHz uplink or 10 MHz + 10 MHz. With augmentation of capacity, such as in 1700/2100 MHz Advanced Wireless Service (AWS) bands, LTE spectrum usage can reach 30 + 30 MHz.

Assumes spectral efficiency of 1.7 bps/Hz in an advanced 4X2 MIMO downlink configuration. Many LTE networks today actually use 2X2 MIMO, which has reduced spectral efficiency of 1.4 bps/Hz. Rysavy Research, *Mobile Broadband Explosion – The 3GPP Wireless Evolution, 4G Americas*, August, 2013.

⁷ Ooyala, *Global Video Index Q1 2014*.

Offloading to Wi-Fi does offer some capacity relief and is the reason so many mobile operators are integrating Wi-Fi into their networks. Unfortunately, Wi-Fi covers only a small percentage of total area, and pressure on cellular capacity will continue unabated.⁸

Because a small handful of heavy users can consume all available network capacity in an area, these heavy users can adversely impact all users in the area. Consequently, policies that slow down certain types of heavy users during periods of congestion can be an important way of ensuring higher-quality service for all customers.

Eventually, 5G technology approaches will dramatically increase capacity, particularly if they can exploit wide new swaths of spectrum in high frequencies and techniques such as massive MIMO, but 5G is only in a research phase at this time, and the hoped-for gains are many years away.

Wireless Network Traffic Management

To achieve high efficiency, wireless networks must perform sophisticated traffic management that involves constant, dynamic changes to connections. Unlike a stable, unvarying wireline connection, the quality of a radio signal can vary moment to moment, especially if a user is mobile, due to complex ways that radio signals propagate, diffract, and reflect in the environment. In addition, because the radio channel is a shared medium, the network must efficiently schedule traffic flow in both directions.


The primary method to accommodate the varying radio signal is to adjust modulation and the amount of error coding. Any modulation method, the alterations made to a waveform to encode digital data, produces a certain error level relative to signal quality (signal-to-noise ratio). With a high-quality signal, the radio can engage a more efficient modulation method, resulting in higher spectral efficiency. In addition, a high-quality signal requires few overhead bits to correct for errors. With a degrading signal, the radio can use more error-correcting bits. Eventually, however, the error rate becomes too high, even with higher error coding; at that point, the radio must change the modulation method to one that is less efficient but able to withstand the noise.

A wireline connection, such as one using a digital subscriber line (DSL) modem, may also have to adjust modulation and coding based on signal quality. Longer distance from a central office is one variable that affects wireline signal quality. The difference relative to wireless, however, is that the connection is far more stable.

⁸ For further discussion of relative roles of Wi-Fi and cellular, see Rysavy Research, “It’s time for a rational perspective on Wi-Fi,” *Gigaom*, April 2007. Available at <http://gigaom.com/2014/04/27/its-time-for-a-rational-perspective-on-wi-fi/>.

Scheduling user communications also improves network efficiency. Figure 2 shows how, at any moment in time, the signal quality for one user can be better than for another. By transmitting to users at times when signal quality is higher, a method called “proportional fair scheduling,” the radio protocols can employ more efficient modulation and less coding overhead. The result is greater network efficiency. This form of efficiency enhancement is a type of traffic discrimination, but one that benefits all users because it increases overall network capacity.

Figure 2: Varying Radio Conditions Between Two Mobile Users



One result of the dynamic modulation and coding used by wireless networks to maximize efficiency is that data rates can vary significantly. For instance, performance at the cell edge is generally poorer than close to the cell tower.

Another variable affecting a user’s experience is the number of users in a coverage area. Unlike fixed networks with fixed endpoints, the number of mobile users can vary, resulting in performance that fluctuates based on the location of all users as well as the types of applications in use. Thus, any requirements for minimum standards for performance, or equivalent performance among users, would wreak havoc on network management—even if such standardization were possible, which it is not.

For the multiple reasons discussed above, users will inevitably experience a wide range of connection quality. The miracle of today’s mobile broadband networks is that they deliver as high a quality-of-experience to as many subscribers as they do. Any regulation that were to interfere with or constrain how this consumer experience is managed, including regulations thwarting necessary traffic prioritization or setting minimum standards of performance, would be detrimental.

Voice, Video, and Quality of Service

Another form of traffic management is prioritizing between voice and data traffic. To offer dependable voice service, wireless networks assign higher priority to voice than to data, true for both circuit-switched voice and future packet-switched voice systems.

Voice over LTE (VoLTE), an operator-provided packet-based voice service seeing growing deployment in 2014, depends on quality-of-service (QoS) mechanisms designed into LTE. The Third Generation Partnership (3GPP) QoS standards were designed by engineers who work closely with industry to determine market needs. The QoS classes in LTE determine packet priority, whether bit rates are guaranteed, the amount of delay through the network, and what percentage of packets can be lost. The highest-priority traffic is traffic that performs control functions for voice and multimedia services. The next highest priority is voice traffic.

LTE will use QoS capabilities, including prioritization, to ensure that packet-based voice meets or exceeds the quality of current circuit-switched voice calls. Without such QoS management, 4G voice would not be possible. A mandate to treat all data equally would have catastrophic consequences for voice users, relegating them to the less spectrally efficient 3G circuit-switched network, which then negatively affects the ability of consumers to reap the benefits of the more advanced 4G networks.

Wireline networks generally handle throughput-sensitive communications, such as voice over IP (VoIP) or video, on a best-effort basis, despite having access to QoS protocols that can prioritize traffic and offer throughput guarantees. Implementing QoS protocols on an end-to-end basis, especially across multiple nodes, is a complex undertaking because it requires negotiation of network resources in advance of application flows. While potentially beneficial to wireline networks in some contexts, most wireline network planners choose to augment capacity rather than implement QoS.

For wireless networks, options to expand capacity are far more limited, so QoS could play a significantly more important role because it will provide additional tools to afford users the best quality of experience. A background firmware update, for example, does not have the same urgency as an interactive video call. Similarly, in that video call, delivering voice data is more important than video because it allows the listener to still understand the speaker even if video momentarily pauses.

Beyond prioritization, the LTE QoS architecture could increase flexibility in service plans by controlling which applications can access the network. Some users, for example, might desire a service plan that allows them unlimited use of all types of content except video.

Table 3 presents examples of applications that could potentially benefit from QoS in LTE. As the industry matures, innovators will find an increasing number of applications that will benefit from QoS control.

Table 3: Application of Quality of Service

Application	How Quality of Service Could be Used
Public Safety	A public-safety network, sharing the same LTE network with consumers, could operate at a high-priority level.
Health Care	Telemedicine applications could benefit from higher-priority video for real-time diagnosis in emergency situations.
Smart Grid	Monitoring and reporting could occur at lower priority, but in a critical situation, such as needing to automatically turn off consumer appliances and AC, messages could travel at higher priority, thus preventing blackouts.
Education	Interactive classrooms could prioritize voice over video. Stored lectures, due to their non-real-time nature, could stream with greater packet delay.

Rules that overlook the unique characteristics and constraints of mobile networks would hobble QoS management, thus degrading the consumer mobile experience as it exists today, removing a potentially useful network management approach that could serve consumers well, and likely eliminating entire classes of applications that haven't even been invented. Who can predict what developers will forgo if they cannot depend on the network? Having to treat all packets and all circumstances as equal in a capacity-constrained environment does not make sense and directly undermines the broad public benefits that mobile broadband is bringing to our economy and our society.

Conclusion

Mobile broadband networks are by their nature and engineering characteristics much more constrained in capacity than wireline networks, mostly due to the finite amount of spectrum and inherent limitations of radio as a medium. All usable radio frequency spectrum, across thousands of licensees, has less than 10 percent of the capacity of a single fiber-optic cable. Furthermore, mobile broadband spectrum represents only a tiny percentage of all usable spectrum.

With rapidly rising consumer mobile broadband usage, as well as the wide proliferation of spectrum-reliant connected devices—the so-called “Internet of things”—U.S. operators are in a continual and accelerating race to stay one step ahead of this demand-versus-capacity challenge. They are doing so by investing over \$30 billion annually to deploy new generations of technology, increasing the number of cell sites, enhancing their use of small cells, and offloading onto Wi-Fi. But their ability to stay out front of this demand wave remains precarious. Success depends on ongoing investment and innovation—both of which hinge on continuing the light regulatory environment that has made the United States a world leader in mobile broadband innovation and deployment.

Wireless technology continues to evolve rapidly, not only with more advanced forms of LTE but also with 5G on the horizon. Development and deployment needs to remain as flexible as possible for wireless to become a competitive alternative to wireline for an increasing number of subscribers.

Mobile operators must manage their networks at multiple layers, including dynamic radio links and appropriate traffic prioritization, along with policies that prevent abuse of the network.

The mobile broadband foundation built on 4G technology is one of mankind's greatest achievements. The success of the wireless broadband sector and its commitment to the maintenance of an open, evolving, and dynamic Internet has not been due to prescriptive government rules, but rather a rare instance of regulatory humility—one that recognizes that while American innovation and consumer behavior can evolve in the blink of an eye, laws and regulations do not. Despite good intentions, additional regulations—particularly given the vibrant state of the U.S. communications sector and the open Internet practices that have evolved in their absence—will likely have adverse, unintended consequences. In particular, undermining consumer-friendly network management and stalling innovation could bring possibly severe economic consequences.

Appendix

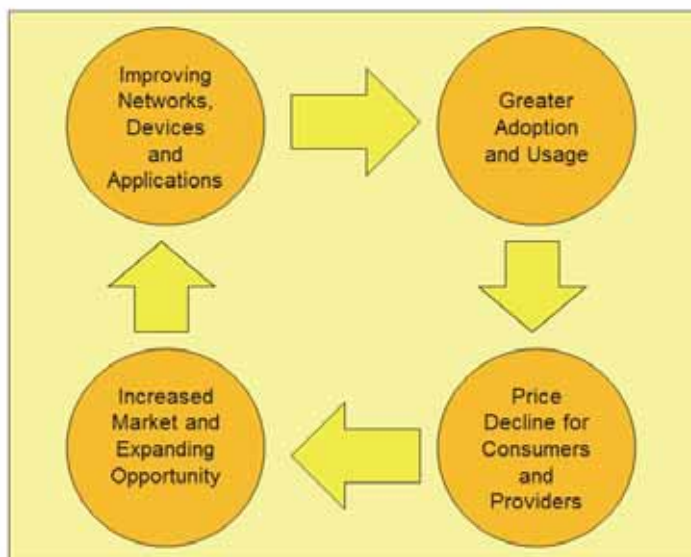
The following are slides used in Rysavy Research presentations to illustrate the profound engineering differences between wireless and wireline networks.

Innovation Areas

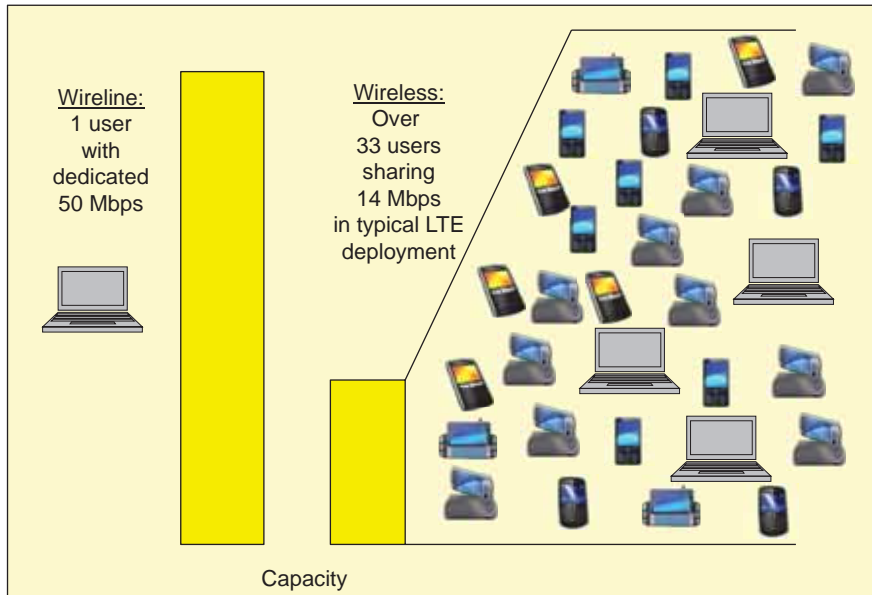
Mobile Broadband: Wireless Technology Innovation Combined with Internet Innovation

Innovation Area	Technologies
Processing	Digital signal processing for radio, huge computing power
Radio	Modulation, multiplexing, antennas
Internet Protocols	IP-based multimedia services including video and voice
Memory	Large capacity storage for songs, photos, video
Mobile Platforms	Smartphones, notebooks, tablets, connected cars, etc.
Applications	Millions
Web Technologies	Internet-hosted applications/content

Virtuous Cycle




Wireline Capacity Compared to Wireless

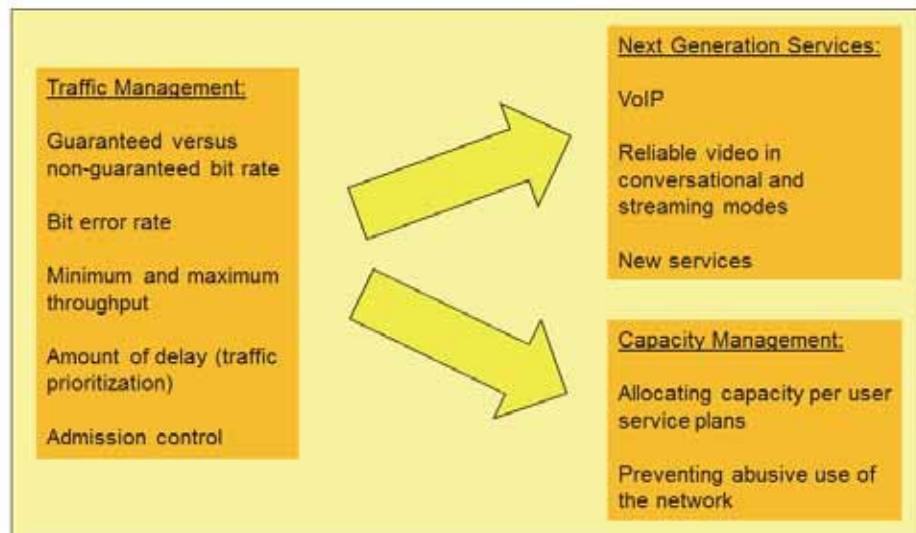


Fiber vs. RF

One fiber strand has more capacity than entire RF spectrum.

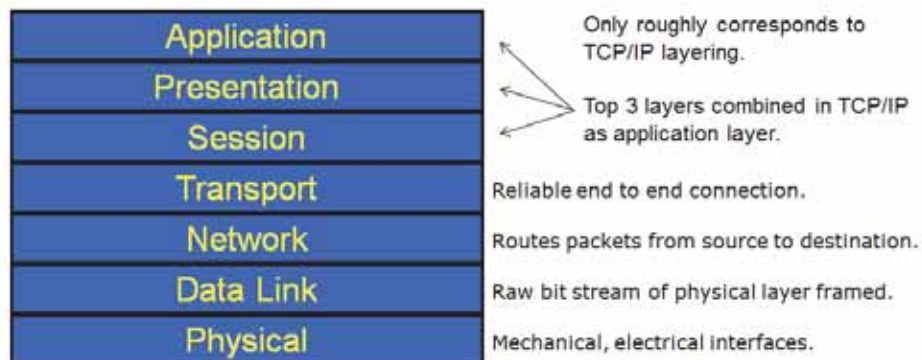


Traffic and Network Management



OSI Network Reference Model

Open Systems Interconnection



Model Details Layer 4 to 7

Layer	Function	Wireline	Wireless
5 to 7	Application	E.g., e-mail, Web	Similar to wireline
4	End-to-end, e.g., TCP/UDP	Most applications use TCP or UDP	TCP/UDP common, TCP fails under congestion, wireless specific transports

Model Details Layer 1 to 3

Layer	Function	Wireline	Wireless
3	Network layer, e.g., IP	All Internet uses IP	Most modern wireless uses IP
2	Link layer for point-to-point and access, varies by network	Many different protocols	Extremely sophisticated methods needed
1	Physical layer, the medium	Wire or fiber	Modulation and coding

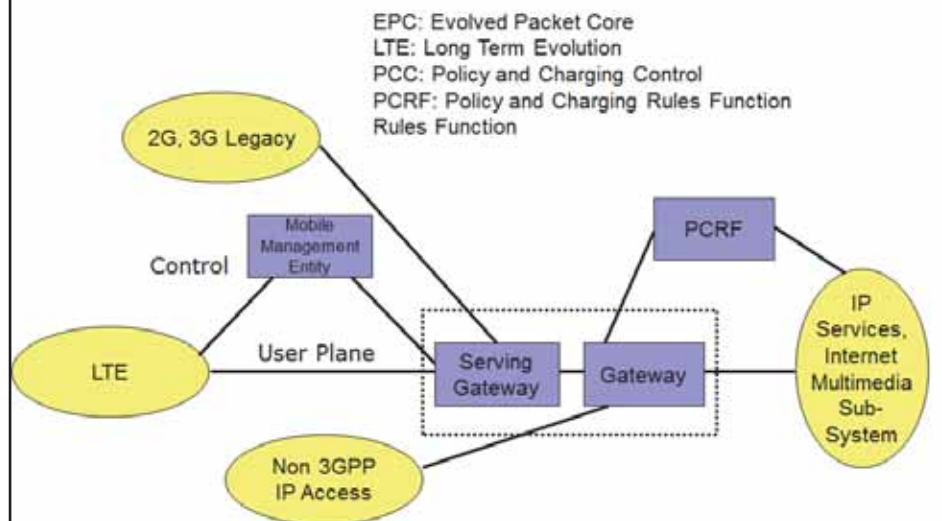
Interaction Between Layers

- Layers supposed to be independent, but ...
- Under congestion:
 - Packet delays, dropped packets.
 - Causes TCP retransmissions and ultimately application failure.
- Much more likely with wireless than wireline.
- Signaling traffic
 - E.g., Session Initiation Protocol (SIP).
 - IP-based but controls traffic streams.

Need to Manage All Layers

Need	Wireless
Congestion Management	Need to deprioritize certain traffic flows. Can be per clearly articulated service plans. Network needs to be cognizant of application type. Congestion less likely to occur on wireline network.
QoS Control and Traffic Management	3GPP option of Policy and Charging Control (PCC). Control of throughputs, delays, priority. Benefits multimedia. Requires application awareness. Essential for very operation of new IP-based core networks. Due to complexity, less likely to be deployed on wireline networks.

QoS-Enabled Networks with Dynamic Policy



Conclusion

- Wireless networks do require specific management approaches at lower layers.
- Helps address interference and achieve high spectral efficiency.
- But network management does not stop at layer two.
- Network management will involve all network layers.
- Wireline networks do not face the same constraints, and unlikely to need these mechanisms.

About Rysavy Research

Peter Rysavy is the president of Rysavy Research LLC, a consulting firm that has specialized in wireless technology since 1993. Projects include analysis of spectrum requirements for mobile broadband, reports on the evolution of wireless technology, evaluation of wireless technology capabilities, strategic consultations, system design, articles, courses and webcasts, network performance measurement, test reports, and acting as expert in patent-litigation cases. Clients include more than seventy-five organizations.

Peter is a broadly published expert on the capabilities and evolution of wireless technology. He has written more than 150 articles, reports, and white papers, and has taught more than forty public wireless courses and webcasts. He has also performed technical evaluations of many wireless technologies including cellular-data services, municipal/mesh Wi-Fi networks, Wi-Fi hotspot networks, mobile browser technologies, wireless e-mail systems, and social networking applications.

From 1988 to 1993, Peter was vice-president of engineering and technology at Traveling Software (later renamed LapLink) where projects included LapLink, LapLink Wireless, and connectivity solutions for a wide variety of mobile platforms. Prior to Traveling Software, he spent seven years at Fluke Corporation where he worked on data-acquisition products and touch-screen technology.

Peter is also the executive director of the Portable Computer and Communications Association (PCCA, <http://www.pcca.org>), an industry organization that evaluates wireless technologies, investigates mobile communications architectures, and promotes wireless-data interoperability. Peter Rysavy graduated with BSEE and MSEE degrees from Stanford University in 1979. More information is available at <http://www.rysavy.com>.



Mobile Future is a coalition of cutting-edge technology and communications companies, consumers and a diverse group of non-profit organizations, working to support an environment that encourages investment and innovation in the dynamic wireless sector. Our mission is to help inform and educate the public and key decision makers in business and government on the broad range of wireless innovations that are transforming our society and the nation's economy.

<http://www.mobilefuture.org>

Copyright © 2014 Rysavy Research, LLC. All rights reserved. <http://www.rysavy.com>.